

✓ Copyright 2018 The TensorFlow Authors.

> Licensed under the Apache License, Version 2.0 (the "License");

[Show code](#)

> MIT License

[Show code](#)

## ✓ Overfit and underfit



[View on TensorFlow.org](#)



[Run in Google Colab](#)



[View source on GitHub](#)



[Download notebook](#)

As always, the code in this example will use the `tf.keras` API, which you can learn more about in the TensorFlow [Keras guide](#).

In both of the previous examples—[classifying text](#) and [predicting fuel efficiency](#)—the accuracy of models on the validation data would peak after training for a number of epochs and then stagnate or start decreasing.

In other words, your model would *overfit* to the training data. Learning how to deal with overfitting is important. Although it's often possible to achieve high accuracy on the *training set*, what you really want is to develop models that generalize well to a *testing set* (or data they haven't seen before).

The opposite of overfitting is *underfitting*. Underfitting occurs when there is still room for improvement on the train data. This can happen for a number of reasons: If the model is not powerful enough, is over-regularized, or has simply not been trained long enough. This means the network has not learned the relevant patterns in the training data.

If you train for too long though, the model will start to overfit and learn patterns from the training data that don't generalize to the test data. You need to strike a balance. Understanding how to train for an appropriate number of epochs as you'll explore below is a useful skill.

To prevent overfitting, the best solution is to use more complete training data. The dataset should cover the full range of inputs that the model is expected to handle. Additional data may only be useful if it covers new and interesting cases.


A model trained on more complete data will naturally generalize better. When that is no longer possible, the next best solution is to use techniques like regularization. These place constraints on the quantity and type of information your model can store. If a network can only afford to memorize a small number of patterns, the optimization process will force it to focus on the most prominent patterns, which have a better chance of generalizing well.

In this notebook, you'll explore several common regularization techniques, and use them to improve on a classification model.

## ✓ Setup

Before getting started, import the necessary packages:

```
1 import tensorflow as tf
2
3 from tensorflow.keras import layers
4 from tensorflow.keras import regularizers
5
6 print(tf.__version__)
```

 2.18.0

```

1 !pip install git+https://github.com/tensorflow/docs
2
3 import tensorflow_docs as tfdocs
4 import tensorflow_docs.modeling
5 import tensorflow_docs.plots

```

```

🔄 Collecting git+https://github.com/tensorflow/docs
  Cloning https://github.com/tensorflow/docs to /tmp/pip-req-build-2rfht4dj
  Running command git clone --filter=blob:none --quiet https://github.com/tensorflow/docs /tmp/pip-req-build-2rfht4dj
  Resolved https://github.com/tensorflow/docs to commit a81ff2e4ebf518b45cad643d88063c6f2b9e5a1e
  Preparing metadata (setup.py) ... done
Collecting astor (from tensorflow-docs==2025.3.6.10029)
  Downloading astor-0.8.1-py2.py3-none-any.whl.metadata (4.2 kB)
Requirement already satisfied: absl-py in /usr/local/lib/python3.11/dist-packages (from tensorflow-docs==2025.3.6.10029)
Requirement already satisfied: Jinja2 in /usr/local/lib/python3.11/dist-packages (from tensorflow-docs==2025.3.6.10029)
Requirement already satisfied: nbformat in /usr/local/lib/python3.11/dist-packages (from tensorflow-docs==2025.3.6.10029)
Requirement already satisfied: protobuf>=3.12 in /usr/local/lib/python3.11/dist-packages (from tensorflow-docs==2025.3.6.10029)
Requirement already satisfied: PyYAML in /usr/local/lib/python3.11/dist-packages (from tensorflow-docs==2025.3.6.10029)
Requirement already satisfied: MarkupSafe>=2.0 in /usr/local/lib/python3.11/dist-packages (from Jinja2->tensorflow-docs==2025.3.6.10029)
Requirement already satisfied: fastjsonschema>=2.15 in /usr/local/lib/python3.11/dist-packages (from nbformat->tensorflow-docs==2025.3.6.10029)
Requirement already satisfied: jsonschema>=2.6 in /usr/local/lib/python3.11/dist-packages (from nbformat->tensorflow-docs==2025.3.6.10029)
Requirement already satisfied: jupyter-core!=5.0.*,>=4.12 in /usr/local/lib/python3.11/dist-packages (from nbformat->tensorflow-docs==2025.3.6.10029)
Requirement already satisfied: traitlets>=5.1 in /usr/local/lib/python3.11/dist-packages (from nbformat->tensorflow-docs==2025.3.6.10029)
Requirement already satisfied: attrs>=22.2.0 in /usr/local/lib/python3.11/dist-packages (from jsonschema>=2.6->nbformat->tensorflow-docs==2025.3.6.10029)
Requirement already satisfied: jsonschema-specifications>=2023.03.6 in /usr/local/lib/python3.11/dist-packages (from jsonschema>=2.6->nbformat->tensorflow-docs==2025.3.6.10029)
Requirement already satisfied: referencing>=0.28.4 in /usr/local/lib/python3.11/dist-packages (from jsonschema>=2.6->nbformat->tensorflow-docs==2025.3.6.10029)
Requirement already satisfied: rpds-py>=0.7.1 in /usr/local/lib/python3.11/dist-packages (from jsonschema>=2.6->nbformat->tensorflow-docs==2025.3.6.10029)
Requirement already satisfied: platformdirs>=2.5 in /usr/local/lib/python3.11/dist-packages (from jupyter-core!=5.0.*,>=4.12->nbformat->tensorflow-docs==2025.3.6.10029)
Requirement already satisfied: typing-extensions>=4.4.0 in /usr/local/lib/python3.11/dist-packages (from referencing>=0.28.4->jsonschema>=2.6->nbformat->tensorflow-docs==2025.3.6.10029)
  Downloading astor-0.8.1-py2.py3-none-any.whl (27 kB)
Building wheels for collected packages: tensorflow-docs
  Building wheel for tensorflow-docs (setup.py) ... done
  Created wheel for tensorflow-docs: filename=tensorflow_docs-2025.3.6.10029-py3-none-any.whl size=186351 sha256=e8709a674
  Stored in directory: /tmp/pip-ephem-wheel-cache-43ih_uds/wheels/34/53/89/3db54cf97ce0f0261aaab3fdc12a847ea0879d34edf373e
Successfully built tensorflow-docs
Installing collected packages: astor, tensorflow-docs
Successfully installed astor-0.8.1 tensorflow-docs-2025.3.6.10029

```

```

1 from IPython import display
2 from matplotlib import pyplot as plt
3
4 import numpy as np
5
6 import pathlib
7 import shutil
8 import tempfile
9
10
11 logdir = pathlib.Path(tempfile.mkdtemp())/"tensorboard_logs"
12 shutil.rmtree(logdir, ignore_errors=True)

```

## ✓ The Higgs dataset

The goal of this tutorial is not to do particle physics, so don't dwell on the details of the dataset. It contains 11,000,000 examples, each with 28 features, and a binary class label.

```
1 gz = tf.keras.utils.get_file('HIGGS.csv.gz', 'http://mlphysics.ics.uci.edu/data/higgs/HIGGS.csv.gz')
```

```

🔄 Downloading data from http://mlphysics.ics.uci.edu/data/higgs/HIGGS.csv.gz
2816407858/2816407858 ————— 290s 0us/step

```

```
1 FEATURES = 28
```

The `tf.data.experimental.CsvDataset` class can be used to read csv records directly from a gzip file with no intermediate decompression step.

```
1 ds = tf.data.experimental.CsvDataset(gz,[float(),]*(FEATURES+1), compression_type="GZIP")
```

That csv reader class returns a list of scalars for each record. The following function repacks that list of scalars into a (feature\_vector, label) pair.

```
1 def pack_row(*row):  
2   label = row[0]  
3   features = tf.stack(row[1:],1)  
4   return features, label
```

TensorFlow is most efficient when operating on large batches of data.

So, instead of repacking each row individually make a new `tf.data.Dataset` that takes batches of 10,000 examples, applies the `pack_row` function to each batch, and then splits the batches back up into individual records:


```
1 packed_ds = ds.batch(10000).map(pack_row).unbatch()
```

Inspect some of the records from this new `packed_ds`.

The features are not perfectly normalized, but this is sufficient for this tutorial.

```
1 for features,label in packed_ds.batch(1000).take(1):  
2   print(features[0])  
3   plt.hist(features.numpy().flatten(), bins = 101)
```

```
tf.Tensor(  
[ 0.8692932 -0.6350818  0.22569026  0.32747006 -0.6899932  0.75420225  
-0.24857314 -1.0920639  0.          1.3749921 -0.6536742  0.9303491  
 1.1074361  1.1389043 -1.5781983 -1.0469854  0.          0.65792954  
-0.01045457 -0.04576717  3.1019614  1.35376  0.9795631  0.97807616  
 0.92000484  0.72165745  0.98875093  0.87667835], shape=(28,), dtype=float32)
```



To keep this tutorial relatively short, use just the first 1,000 samples for validation, and the next 10,000 for training:

```
1 N_VALIDATION = int(1e3)  
2 N_TRAIN = int(1e4)  
3 BUFFER_SIZE = int(1e4)  
4 BATCH_SIZE = 500  
5 STEPS_PER_EPOCH = N_TRAIN//BATCH_SIZE
```

The `Dataset.skip` and `Dataset.take` methods make this easy.

At the same time, use the `Dataset.cache` method to ensure that the loader doesn't need to re-read the data from the file on each epoch:

```
1 validate_ds = packed_ds.take(N_VALIDATION).cache()
2 train_ds = packed_ds.skip(N_VALIDATION).take(N_TRAIN).cache()

1 train_ds
↳ <CacheDataset element_spec=(TensorSpec(shape=(28,), dtype=tf.float32, name=None), TensorSpec(shape=(), dtype=tf.float32, name=None))>
```

These datasets return individual examples. Use the `Dataset.batch` method to create batches of an appropriate size for training. Before batching, also remember to use `Dataset.shuffle` and `Dataset.repeat` on the training set.

```
1 validate_ds = validate_ds.batch(BATCH_SIZE)
2 train_ds = train_ds.shuffle(BUFFER_SIZE).repeat().batch(BATCH_SIZE)
```

## ✓ Demonstrate overfitting

The simplest way to prevent overfitting is to start with a small model: A model with a small number of learnable parameters (which is determined by the number of layers and the number of units per layer). In deep learning, the number of learnable parameters in a model is often referred to as the model's "capacity".

Intuitively, a model with more parameters will have more "memorization capacity" and therefore will be able to easily learn a perfect dictionary-like mapping between training samples and their targets, a mapping without any generalization power, but this would be useless when making predictions on previously unseen data.

Always keep this in mind: deep learning models tend to be good at fitting to the training data, but the real challenge is generalization, not fitting.

On the other hand, if the network has limited memorization resources, it will not be able to learn the mapping as easily. To minimize its loss, it will have to learn compressed representations that have more predictive power. At the same time, if you make your model too small, it will have difficulty fitting to the training data. There is a balance between "too much capacity" and "not enough capacity".

Unfortunately, there is no magical formula to determine the right size or architecture of your model (in terms of the number of layers, or the right size for each layer). You will have to experiment using a series of different architectures.

To find an appropriate model size, it's best to start with relatively few layers and parameters, then begin increasing the size of the layers or adding new layers until you see diminishing returns on the validation loss.

Start with a simple model using only densely-connected layers (`tf.keras.layers.Dense`) as a baseline, then create larger models, and compare them.


## ✓ Training procedure

Many models train better if you gradually reduce the learning rate during training. Use `tf.keras.optimizers.schedules` to reduce the learning rate over time:

```
1 lr_schedule = tf.keras.optimizers.schedules.InverseTimeDecay(
2     0.001,
3     decay_steps=STEPS_PER_EPOCH*1000,
4     decay_rate=1,
5     staircase=False)
6
7 def get_optimizer():
8     return tf.keras.optimizers.Adam(lr_schedule)
```

The code above sets a `tf.keras.optimizers.schedules.InverseTimeDecay` to hyperbolically decrease the learning rate to 1/2 of the base rate at 1,000 epochs, 1/3 at 2,000 epochs, and so on.

```
1 step = np.linspace(0,100000)
2 lr = lr_schedule(step)
3 plt.figure(figsize = (8,6))
4 plt.plot(step/STEPS_PER_EPOCH, lr)
5 plt.ylim([0,max(plt.ylim())])
6 plt.xlabel('Epoch')
7 _ = plt.ylabel('Learning Rate')
8
```



Each model in this tutorial will use the same training configuration. So set these up in a reusable way, starting with the list of callbacks.

The training for this tutorial runs for many short epochs. To reduce the logging noise use the `tfdocs.EpochDots` which simply prints a `.` for each epoch, and a full set of metrics every 100 epochs.

Next include `tf.keras.callbacks.EarlyStopping` to avoid long and unnecessary training times. Note that this callback is set to monitor the `val_binary_crossentropy`, not the `val_loss`. This difference will be important later.

Use `callbacks.TensorBoard` to generate TensorBoard logs for the training.

```
1 def get_callbacks(name):
2     return [
3         tfdocs.modeling.EpochDots(),
4         tf.keras.callbacks.EarlyStopping(monitor='val_binary_crossentropy', patience=200),
5         tf.keras.callbacks.TensorBoard(logdir/name),
6     ]
```

Similarly each model will use the same `Model.compile` and `Model.fit` settings:

```
1 def compile_and_fit(model, name, optimizer=None, max_epochs=10000):
2     if optimizer is None:
3         optimizer = get_optimizer()
4     model.compile(optimizer=optimizer,
5                 loss=tf.keras.losses.BinaryCrossentropy(from_logits=True),
```

```

6         metrics=[
7             tf.keras.metrics.BinaryCrossentropy(
8                 from_logits=True, name='binary_crossentropy'),
9             'accuracy'])
10
11 model.summary()
12
13 history = model.fit(
14     train_ds.map(lambda x, y: (x, tf.expand_dims(y, axis=-1))),
15     steps_per_epoch = STEPS_PER_EPOCH,
16     epochs=max_epochs,
17     validation_data=validate_ds.map(lambda x, y: (x, tf.expand_dims(y, axis=-1))),
18     callbacks=get_callbacks(name),
19     verbose=0)
20 return history

```

## ✓ Tiny model

Start by training a model:

```

1 tiny_model = tf.keras.Sequential([
2     layers.Dense(16, activation='elu', input_shape=(FEATURES,)),
3     layers.Dense(1)
4 ])

```

```

➔ /usr/local/lib/python3.11/dist-packages/keras/src/layers/core/dense.py:87: UserWarning: Do not pass an `input_shape`/`input_shape_tuple` argument to `Dense` layer.
super().__init__(activity_regularizer=activity_regularizer, **kwargs)

```

```
1 size_histories = {}
```

```
1 size_histories['Tiny'] = compile_and_fit(tiny_model, 'sizes/Tiny')
```

```
➔ Model: "sequential"
```

Layer (type)	Output Shape	Param #
dense (Dense)	(None, 16)	464
dense_1 (Dense)	(None, 1)	17

```

Total params: 481 (1.88 KB)
Trainable params: 481 (1.88 KB)
Non-trainable params: 0 (0.00 B)

```

```


Epoch: 0, accuracy:0.5069, binary_crossentropy:0.8599, loss:0.8599, val_accuracy:0.5020, val_binary_crossentropy:0.811
./usr/local/lib/python3.11/dist-packages/keras/src/trainers/epoch_iterator.py:151: UserWarning: Your input ran out of data:
self._interrupted_warning()
.....
Epoch: 100, accuracy:0.6051, binary_crossentropy:0.6266, loss:0.6266, val_accuracy:0.5800, val_binary_crossentropy:0.611
.....
Epoch: 200, accuracy:0.6273, binary_crossentropy:0.6128, loss:0.6128, val_accuracy:0.5970, val_binary_crossentropy:0.611
.....
Epoch: 300, accuracy:0.6379, binary_crossentropy:0.6033, loss:0.6033, val_accuracy:0.6240, val_binary_crossentropy:0.611
.....
Epoch: 400, accuracy:0.6582, binary_crossentropy:0.5960, loss:0.5960, val_accuracy:0.6300, val_binary_crossentropy:0.611
.....
Epoch: 500, accuracy:0.6559, binary_crossentropy:0.5908, loss:0.5908, val_accuracy:0.6400, val_binary_crossentropy:0.611
.....
Epoch: 600, accuracy:0.6590, binary_crossentropy:0.5872, loss:0.5872, val_accuracy:0.6430, val_binary_crossentropy:0.611
.....
Epoch: 700, accuracy:0.6675, binary_crossentropy:0.5844, loss:0.5844, val_accuracy:0.6400, val_binary_crossentropy:0.611
.....
Epoch: 800, accuracy:0.6639, binary_crossentropy:0.5830, loss:0.5830, val_accuracy:0.6330, val_binary_crossentropy:0.611
.....

```

Now check how the model did:

```
1 plotter = tfdocs.plots.HistoryPlotter(metric = 'binary_crossentropy', smoothing_std=10)
2 plotter.plot(size_histories)
3 plt.ylim([0.5, 0.7])
```

↔ (0.5, 0.7)



## ✓ Small model

To check if you can beat the performance of the small model, progressively train some larger models.

Try two hidden layers with 16 units each:

```
1 small_model = tf.keras.Sequential([
2     # `input_shape` is only required here so that `.summary` works.
3     layers.Dense(16, activation='elu', input_shape=(FEATURES,)),
4     layers.Dense(16, activation='elu'),
5     layers.Dense(1)
6 ])

1 size_histories['Small'] = compile_and_fit(small_model, 'sizes/Small')
```

Model: "sequential\_1"

Layer (type)	Output Shape	Param #
dense_2 (Dense)	(None, 16)	464
dense_3 (Dense)	(None, 16)	272
dense_4 (Dense)	(None, 1)	17

Total params: 753 (2.94 KB)  
Trainable params: 753 (2.94 KB)  
Non-trainable params: 0 (0.00 B)

```
Epoch: 0, accuracy:0.4751, binary_crossentropy:0.7112, loss:0.7112, val_accuracy:0.4580, val_binary_crossentropy:0.69!
.....
Epoch: 100, accuracy:0.6290, binary_crossentropy:0.6104, loss:0.6104, val_accuracy:0.6300, val_binary_crossentropy:0.6!
.....
Epoch: 200, accuracy:0.6619, binary_crossentropy:0.5857, loss:0.5857, val_accuracy:0.6730, val_binary_crossentropy:0.6!
.....
Epoch: 300, accuracy:0.6731, binary_crossentropy:0.5762, loss:0.5762, val_accuracy:0.6610, val_binary_crossentropy:0.6!
.....
Epoch: 400, accuracy:0.6823, binary_crossentropy:0.5691, loss:0.5691, val_accuracy:0.6740, val_binary_crossentropy:0.6!
.....
Epoch: 500, accuracy:0.6819, binary_crossentropy:0.5654, loss:0.5654, val_accuracy:0.6530, val_binary_crossentropy:0.6!
.....
Epoch: 600, accuracy:0.6905, binary_crossentropy:0.5617, loss:0.5617, val_accuracy:0.6690, val_binary_crossentropy:0.6!
.....
Epoch: 700, accuracy:0.6938, binary_crossentropy:0.5590, loss:0.5590, val_accuracy:0.6720, val_binary_crossentropy:0.6!
.....
Epoch: 800, accuracy:0.6960, binary_crossentropy:0.5556, loss:0.5556, val_accuracy:0.6790, val_binary_crossentropy:0.6!
.....
Epoch: 900, accuracy:0.6972, binary_crossentropy:0.5525, loss:0.5525, val_accuracy:0.6700, val_binary_crossentropy:0.6!
.....
Epoch: 1000, accuracy:0.7018, binary_crossentropy:0.5500, loss:0.5500, val_accuracy:0.6660, val_binary_crossentropy:0!
.....
```

## Medium model

Now try three hidden layers with 64 units each:

```
1 medium_model = tf.keras.Sequential([
2     layers.Dense(64, activation='elu', input_shape=(FEATURES,)),
3     layers.Dense(64, activation='elu'),
4     layers.Dense(64, activation='elu'),
5     layers.Dense(1)
6 ])
```

And train the model using the same data:

```
1 size_histories['Medium'] = compile_and_fit(medium_model, "sizes/Medium")
```

Model: "sequential\_2"

Layer (type)	Output Shape	Param #
dense_5 (Dense)	(None, 64)	1,856
dense_6 (Dense)	(None, 64)	4,160
dense_7 (Dense)	(None, 64)	4,160
dense_8 (Dense)	(None, 1)	65

Total params: 10,241 (40.00 KB)  
Trainable params: 10,241 (40.00 KB)  
Non-trainable params: 0 (0.00 B)

Epoch: 0, accuracy:0.4952, binary\_crossentropy:0.6848, loss:0.6848, val\_accuracy:0.5030, val\_binary\_crossentropy:0.67  
Epoch: 100, accuracy:0.7148, binary\_crossentropy:0.5259, loss:0.5259, val\_accuracy:0.6720, val\_binary\_crossentropy:0.6  
Epoch: 200, accuracy:0.7875, binary\_crossentropy:0.4266, loss:0.4266, val\_accuracy:0.6460, val\_binary\_crossentropy:0.6

## Large model

As an exercise, you can create an even larger model and check how quickly it begins overfitting. Next, add to this benchmark a network that has much more capacity, far more than the problem would warrant:

```
1 large_model = tf.keras.Sequential([  
2     layers.Dense(512, activation='elu', input_shape=(FEATURES,)),  
3     layers.Dense(512, activation='elu'),  
4     layers.Dense(512, activation='elu'),  
5     layers.Dense(512, activation='elu'),  
6     layers.Dense(1)  
7 ])
```

And, again, train the model using the same data:

```
1 size_histories['large'] = compile_and_fit(large_model, "sizes/large")
```

Model: "sequential\_3"

Layer (type)	Output Shape	Param #
dense_9 (Dense)	(None, 512)	14,848
dense_10 (Dense)	(None, 512)	262,656
dense_11 (Dense)	(None, 512)	262,656
dense_12 (Dense)	(None, 512)	262,656
dense_13 (Dense)	(None, 1)	513

Total params: 803,329 (3.06 MB)  
Trainable params: 803,329 (3.06 MB)  
Non-trainable params: 0 (0.00 B)

Epoch: 0, accuracy:0.5085, binary\_crossentropy:0.8122, loss:0.8122, val\_accuracy:0.4780, val\_binary\_crossentropy:0.68  
Epoch: 100, accuracy:1.0000, binary\_crossentropy:0.0020, loss:0.0020, val\_accuracy:0.6330, val\_binary\_crossentropy:1.6  
Epoch: 200, accuracy:1.0000, binary\_crossentropy:0.0001, loss:0.0001, val\_accuracy:0.6390, val\_binary\_crossentropy:2.4

## Plot the training and validation losses

The solid lines show the training loss, and the dashed lines show the validation loss (remember: a lower validation loss indicates a better model).

While building a larger model gives it more power, if this power is not constrained somehow it can easily overfit to the training set.

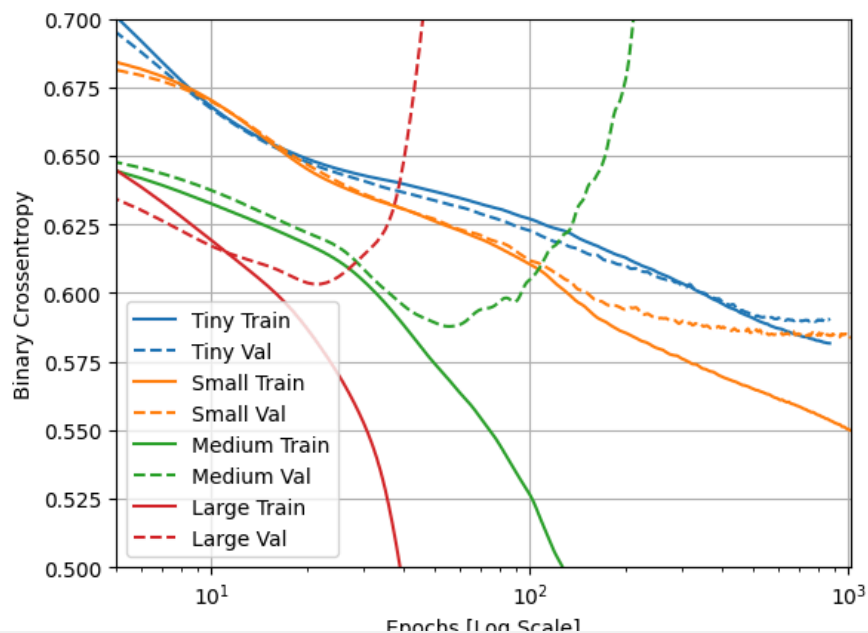
In this example, typically, only the "Tiny" model manages to avoid overfitting altogether, and each of the larger models overfit the data more quickly. This becomes so severe for the "large" model that you need to switch the plot to a log-scale to really figure out what's happening.

This is apparent if you plot and compare the validation metrics to the training metrics.

- It's normal for there to be a small difference.
- If both metrics are moving in the same direction, everything is fine.
- If the validation metric begins to stagnate while the training metric continues to improve, you are probably close to overfitting.
- If the validation metric is going in the wrong direction, the model is clearly overfitting.

```
1 plotter.plot(size_histories)
2 a = plt.xscale('log')
3 plt.xlim([5, max(plt.xlim())])
4 plt.ylim([0.5, 0.7])
5 plt.xlabel("Epochs [Log Scale]")
```

Text(0.5, 0, 'Epochs [Log Scale]')



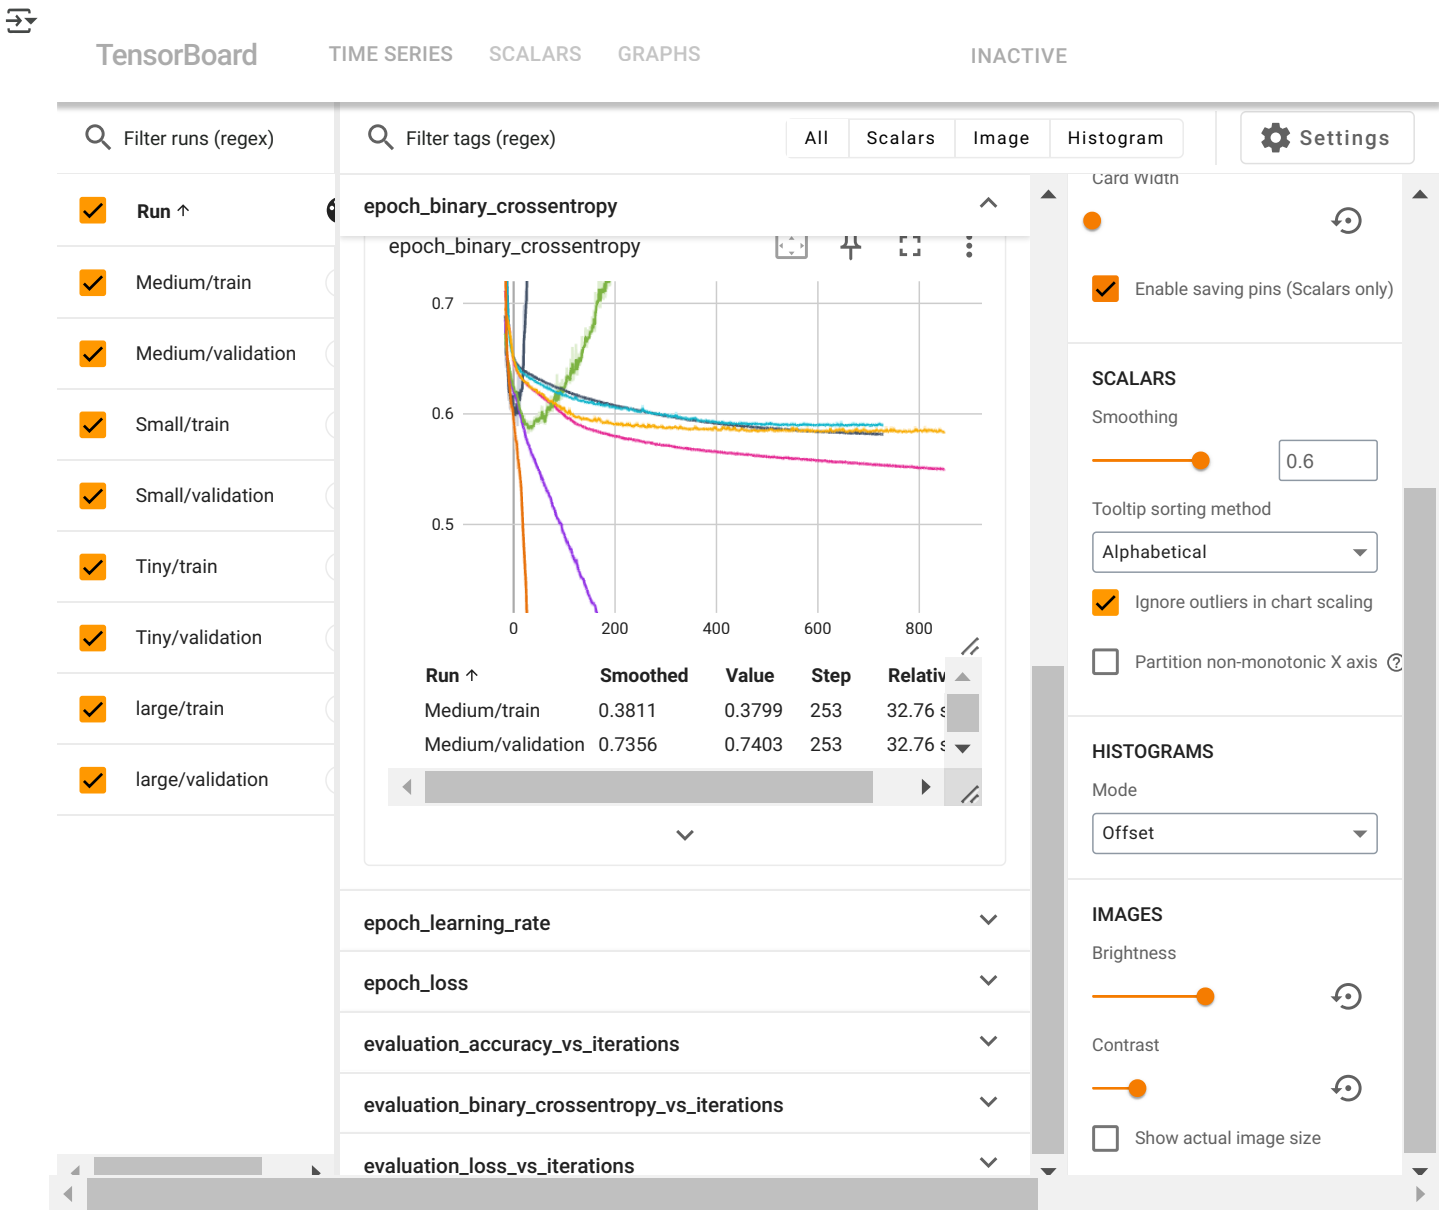
Note: All the above training runs used the `callbacks.EarlyStopping` to end the training once it was clear the model was not making progress.

## View in TensorBoard

These models all wrote TensorBoard logs during training.

Open an embedded TensorBoard viewer inside a notebook (Sorry, this doesn't display on tensorflow.org):

```
1 # Load the TensorBoard notebook extension
2 %load_ext tensorboard
3
4 # Open an embedded TensorBoard viewer
5 %tensorboard --logdir {logdir}/sizes
```



You can view the [results of a previous run](#) of this notebook on [TensorBoard.dev](#).

## ✓ Strategies to prevent overfitting

Before getting into the content of this section copy the training logs from the "Tiny" model above, to use as a baseline for comparison.

```

1 shutil.rmtree(logdir/'regularizers/Tiny', ignore_errors=True)
2 shutil.copytree(logdir/'sizes/Tiny', logdir/'regularizers/Tiny')

🔗 PosixPath('/tmp/tmpb3_r55c2/tensorboard_logs/regularizers/Tiny')

1 regularizer_histories = {}
2 regularizer_histories['Tiny'] = size_histories['Tiny']

```

## ✓ Add weight regularization

You may be familiar with Occam's Razor principle: given two explanations for something, the explanation most likely to be correct is the "simplest" one, the one that makes the least amount of assumptions. This also applies to the models learned by neural networks: given

some training data and a network architecture, there are multiple sets of weights values (multiple models) that could explain the data, and simpler models are less likely to overfit than complex ones.

A "simple model" in this context is a model where the distribution of parameter values has less entropy (or a model with fewer parameters altogether, as demonstrated in the section above). Thus a common way to mitigate overfitting is to put constraints on the complexity of a network by forcing its weights only to take small values, which makes the distribution of weight values more "regular". This is called "weight regularization", and it is done by adding to the loss function of the network a cost associated with having large weights. This cost comes in two flavors:

- [L1 regularization](#), where the cost added is proportional to the absolute value of the weights coefficients (i.e. to what is called the "L1 norm" of the weights).
- [L2 regularization](#), where the cost added is proportional to the square of the value of the weights coefficients (i.e. to what is called the squared "L2 norm" of the weights). L2 regularization is also called weight decay in the context of neural networks. Don't let the different name confuse you: weight decay is mathematically the exact same as L2 regularization.

L1 regularization pushes weights towards exactly zero, encouraging a sparse model. L2 regularization will penalize the weights parameters without making them sparse since the penalty goes to zero for small weights—one reason why L2 is more common.

In `tf.keras`, weight regularization is added by passing weight regularizer instances to layers as keyword arguments. Add L2 weight regularization:

```
1 l2_model = tf.keras.Sequential([
2     layers.Dense(512, activation='elu',
3         kernel_regularizer=regularizers.l2(0.001),
4         input_shape=(FEATURES,)),
5     layers.Dense(512, activation='elu',
6         kernel_regularizer=regularizers.l2(0.001)),
7     layers.Dense(512, activation='elu',
8         kernel_regularizer=regularizers.l2(0.001)),
9     layers.Dense(512, activation='elu',
10        kernel_regularizer=regularizers.l2(0.001)),
11    layers.Dense(1)
12 ])
13
14 regularizer_histories['l2'] = compile_and_fit(l2_model, "regularizers/l2")
```

```

/usr/local/lib/python3.11/dist-packages/keras/src/layers/core/dense.py:87: UserWarning: Do not pass an `input_shape`/`input_shape_tuple` argument to the `Dense` layer constructor.
super().__init__(activity_regularizer=activity_regularizer, **kwargs)

```

Model: "sequential\_4"

Layer (type)	Output Shape	Param #
dense_14 (Dense)	(None, 512)	14,848
dense_15 (Dense)	(None, 512)	262,656
dense_16 (Dense)	(None, 512)	262,656
dense_17 (Dense)	(None, 512)	262,656
dense_18 (Dense)	(None, 1)	513

Total params: 803,329 (3.06 MB)  
 Trainable params: 803,329 (3.06 MB)  
 Non-trainable params: 0 (0.00 B)

```

Epoch: 0, accuracy:0.5009, binary_crossentropy:0.7780, loss:2.2853, val_accuracy:0.4670, val_binary_crossentropy:0.684
./usr/local/lib/python3.11/dist-packages/keras/src/trainers/epoch_iterator.py:151: UserWarning: Your input ran out of data:
self._interrupted_warning()

```

```

.....
Epoch: 100, accuracy:0.6606, binary_crossentropy:0.5928, loss:0.6168, val_accuracy:0.6620, val_binary_crossentropy:0.61
.....
Epoch: 200, accuracy:0.6755, binary_crossentropy:0.5822, loss:0.6060, val_accuracy:0.6680, val_binary_crossentropy:0.61
.....
Epoch: 300, accuracy:0.6781, binary_crossentropy:0.5747, loss:0.5981, val_accuracy:0.6610, val_binary_crossentropy:0.61
.....
Epoch: 400, accuracy:0.6918, binary_crossentropy:0.5653, loss:0.5897, val_accuracy:0.6730, val_binary_crossentropy:0.61
.....
Epoch: 500, accuracy:0.6888, binary_crossentropy:0.5616, loss:0.5904, val_accuracy:0.6710, val_binary_crossentropy:0.61
.....

```

12(0.001) means that every coefficient in the weight matrix of the layer will add  $0.001 * \text{weight\_coefficient\_value}^2$  to the total **loss** of the network.

That is why you need to monitor the `binary_crossentropy` directly. Because it doesn't have this regularization component mixed in.


So, that same "Large" model with an L2 regularization penalty performs much better:

```

1 plotter.plot(regularizer_histories)
2 plt.ylim([0.5, 0.7])

```

(0.5, 0.7)



As demonstrated in the diagram above, the "L2" regularized model is now much more competitive with the "Tiny" model. This "L2" model is also much more resistant to overfitting than the "Large" model it was based on despite having the same number of parameters.

## ✓ More info

There are two important things to note about this sort of regularization:

1. If you are writing your own training loop, then you need to be sure to ask the model for its regularization losses.

```
1 result = l2_model(features)
2 regularization_loss=tf.add_n(l2_model.losses)
```

2. This implementation works by adding the weight penalties to the model's loss, and then applying a standard optimization procedure after that.

There is a second approach that instead only runs the optimizer on the raw loss, and then while applying the calculated step the optimizer also applies some weight decay. This "decoupled weight decay" is used in optimizers like `tf.keras.optimizers.Ftrl` and `tfa.optimizers.AdamW`.

## ✓ Add dropout

Dropout is one of the most effective and most commonly used regularization techniques for neural networks, developed by Hinton and his students at the University of Toronto.

The intuitive explanation for dropout is that because individual nodes in the network cannot rely on the output of the others, each node must output features that are useful on their own.

Dropout, applied to a layer, consists of randomly "dropping out" (i.e. set to zero) a number of output features of the layer during training. For example, a given layer would normally have returned a vector `[0.2, 0.5, 1.3, 0.8, 1.1]` for a given input sample during training; after applying dropout, this vector will have a few zero entries distributed at random, e.g. `[0, 0.5, 1.3, 0, 1.1]`.

The "dropout rate" is the fraction of the features that are being zeroed-out; it is usually set between 0.2 and 0.5. At test time, no units are dropped out, and instead the layer's output values are scaled down by a factor equal to the dropout rate, so as to balance for the fact that more units are active than at training time.

In Keras, you can introduce dropout in a network via the `tf.keras.layers.Dropout` layer, which gets applied to the output of layer right before.

Add two dropout layers to your network to check how well they do at reducing overfitting:

```
1 dropout_model = tf.keras.Sequential([
2     layers.Dense(512, activation='elu', input_shape=(FEATURES,)),
3     layers.Dropout(0.5),
4     layers.Dense(512, activation='elu'),
5     layers.Dropout(0.5),
6     layers.Dense(512, activation='elu'),
7     layers.Dropout(0.5),
8     layers.Dense(512, activation='elu'),
9     layers.Dropout(0.5),
10    layers.Dense(1)
11 ])
12
13 regularizer_histories['dropout'] = compile_and_fit(dropout_model, "regularizers/dropout")
```

Model: "sequential\_5"

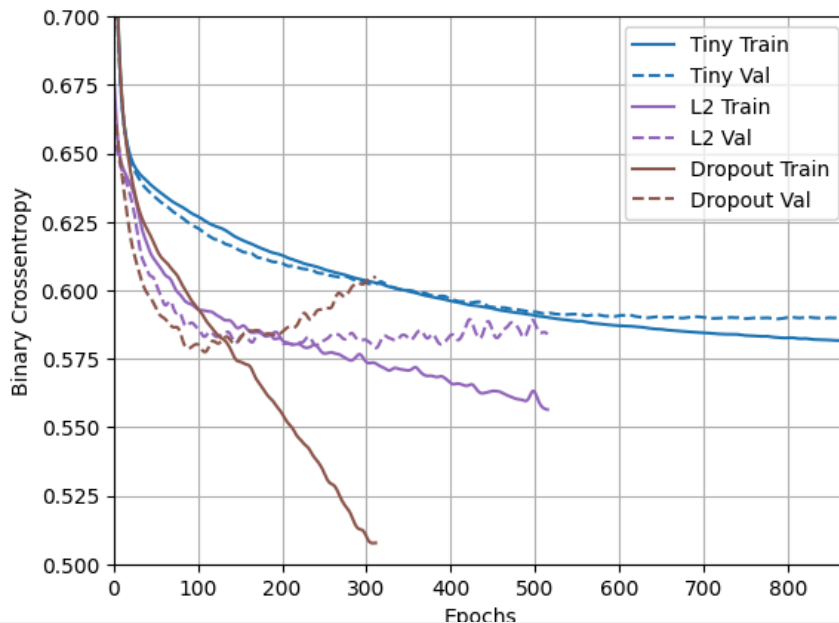
Layer (type)	Output Shape	Param #
dense_19 (Dense)	(None, 512)	14,848
dropout (Dropout)	(None, 512)	0
dense_20 (Dense)	(None, 512)	262,656
dropout_1 (Dropout)	(None, 512)	0
dense_21 (Dense)	(None, 512)	262,656
dropout_2 (Dropout)	(None, 512)	0
dense_22 (Dense)	(None, 512)	262,656
dropout_3 (Dropout)	(None, 512)	0
dense_23 (Dense)	(None, 1)	513

Total params: 803,329 (3.06 MB)  
Trainable params: 803,329 (3.06 MB)  
Non-trainable params: 0 (0.00 B)

```
Epoch: 0, accuracy:0.4970, binary_crossentropy:0.8047, loss:0.8047, val_accuracy:0.4620, val_binary_crossentropy:0.6847
Epoch: 100, accuracy:0.6566, binary_crossentropy:0.5953, loss:0.5953, val_accuracy:0.6640, val_binary_crossentropy:0.6047
Epoch: 200, accuracy:0.6895, binary_crossentropy:0.5574, loss:0.5574, val_accuracy:0.6790, val_binary_crossentropy:0.5947
Epoch: 300, accuracy:0.7264, binary_crossentropy:0.5143, loss:0.5143, val_accuracy:0.6940, val_binary_crossentropy:0.5847
```

```
1 plotter.plot(regularizer_histories)
2 plt.ylim([0.5, 0.7])
```

(0.5, 0.7)



It's clear from this plot that both of these regularization approaches improve the behavior of the "Large" model. But this still doesn't beat even the "Tiny" baseline.

Next try them both, together, and see if that does better.

## ✓ Combined L2 + dropout

```

1 combined_model = tf.keras.Sequential([
2     layers.Dense(512, kernel_regularizer=regularizers.l2(0.0001),
3         activation='elu', input_shape=(FEATURES,)),
4     layers.Dropout(0.5),
5     layers.Dense(512, kernel_regularizer=regularizers.l2(0.0001),
6         activation='elu'),
7     layers.Dropout(0.5),
8     layers.Dense(512, kernel_regularizer=regularizers.l2(0.0001),
9         activation='elu'),
10    layers.Dropout(0.5),
11    layers.Dense(512, kernel_regularizer=regularizers.l2(0.0001),
12        activation='elu'),
13    layers.Dropout(0.5),
14    layers.Dense(1)
15 ])
16
17 regularizer_histories['combined'] = compile_and_fit(combined_model, "regularizers/combined")

```

→ Model: "sequential\_6"

Layer (type)	Output Shape	Param #
dense_24 (Dense)	(None, 512)	14,848
dropout_4 (Dropout)	(None, 512)	0
dense_25 (Dense)	(None, 512)	262,656
dropout_5 (Dropout)	(None, 512)	0
dense_26 (Dense)	(None, 512)	262,656
dropout_6 (Dropout)	(None, 512)	0
dense_27 (Dense)	(None, 512)	262,656
dropout_7 (Dropout)	(None, 512)	0
dense_28 (Dense)	(None, 1)	513

**Total params:** 803,329 (3.06 MB)  
**Trainable params:** 803,329 (3.06 MB)  
**Non-trainable params:** 0 (0.00 B)

```

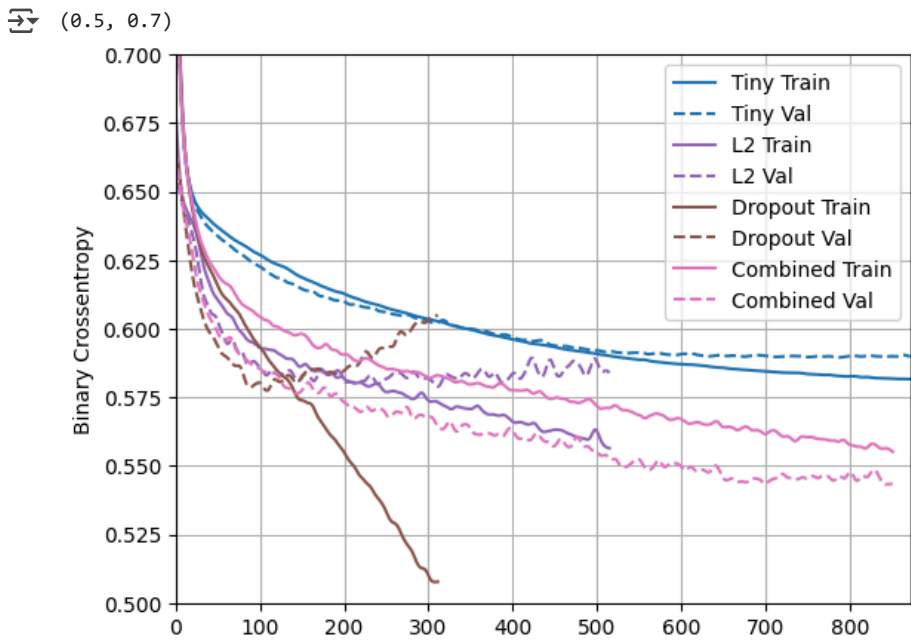
Epoch: 0, accuracy:0.5017, binary_crossentropy:0.7989, loss:0.9572, val_accuracy:0.5170, val_binary_crossentropy:0.68:
.....
Epoch: 100, accuracy:0.6496, binary_crossentropy:0.6078, loss:0.6376, val_accuracy:0.6320, val_binary_crossentropy:0.!:
.....
Epoch: 200, accuracy:0.6611, binary_crossentropy:0.5926, loss:0.6179, val_accuracy:0.6860, val_binary_crossentropy:0.!:
.....
Epoch: 300, accuracy:0.6682, binary_crossentropy:0.5833, loss:0.6114, val_accuracy:0.6730, val_binary_crossentropy:0.!:
.....
Epoch: 400, accuracy:0.6769, binary_crossentropy:0.5770, loss:0.6066, val_accuracy:0.6990, val_binary_crossentropy:0.!:
.....
Epoch: 500, accuracy:0.6811, binary_crossentropy:0.5707, loss:0.6021, val_accuracy:0.6890, val_binary_crossentropy:0.!:
.....
Epoch: 600, accuracy:0.6810, binary_crossentropy:0.5681, loss:0.6017, val_accuracy:0.7060, val_binary_crossentropy:0.!:
.....
Epoch: 700, accuracy:0.6904, binary_crossentropy:0.5578, loss:0.5932, val_accuracy:0.6930, val_binary_crossentropy:0.!:
.....
Epoch: 800, accuracy:0.6866, binary_crossentropy:0.5575, loss:0.5945, val_accuracy:0.7060, val_binary_crossentropy:0.!:
.....

```

```

1 plotter.plot(regularizer_histories)
2 plt.ylim([0.5, 0.7])

```



This model with the "Combined" regularization is obviously the best one so far.

View in TensorBoard

These models also recorded TensorBoard logs.

To open an embedded run the following into a code-cell (Sorry, this doesn't display on tensorflow.org):

```
1 %tensorboard --logdir {logdir}/regularizers
```

